Applied artificial intelligence

by Jerry Felsen, Ph.D.

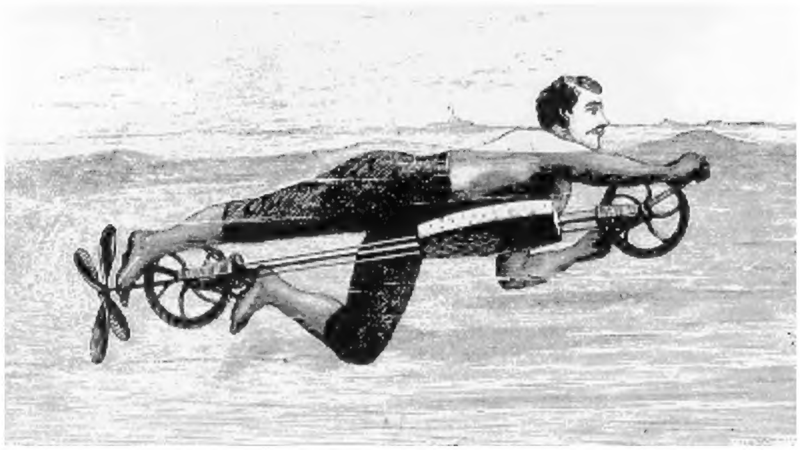

Both the science of cybernetics and artificial intelligence (AI) techniques are central to the design of viable man-computer investment decision systems.

Cybernetics has been defined as the interdisciplinary science of communication and control in the animal and machine. Its major province is the study of systems which may be exceedingly complex and dynamic . . . perhaps defying a full description. These are exactly the sort of systems faced by the investment analyst.

The cybernetic approach to investment analysis was not feasible before the 1970s because there were no suitable mathematical techniques for efficiently processing the extremely complex information patterns characterizing the stock market. Artificial intelligence techniques, especially pattern recognition and machine learning, now enable us to efficiently synthesize large amounts of relevant information into investment decisions. This enables us to formulate new investment techniques which do not have many of the limitations of conventional methods.

A cybernetician regards the stock market as a complex, viable, dynamic, and adaptive system. It can change its internal structure in response to changes in its environment. It maintains multiple contacts with its environment, and it can adapt to unforeseen circumstances. A decision system that operates in the stock market environment, in order to be successful, must possess similarly viable characteristics. That is, it must be able to learn and to adapt to changes in the market environment.

Complexity required

A basic cybernetic principle, the law of requisite variety (1), requires that the complexity of the analysis must approach the variety of the system to be analyzed if analysis is to be successful. This means that, ideally, the investment decision system must consider all relevant factors affecting the market. This concept is very important. Variety in the real world is to be handled by an equivalent variety in the decision system and cannot be competently handled by less.

We frequently mislead ourselves into thinking that we can outwit the natural law of requisite variety in

VISIT…

LANZAROTE
Caliente.COM

the stock market, just as many imagine that they can beat other natural laws at the racetrack or at the casino. This means that complex, high-variety decision problems cannot be solved effectively by simple, low variety methods, such as the usual technical indicators.

Technical analysis

Frequently, in contemporary technical analysis, investment decisions are based only on a small segment of the information spectrum characterizing the market. In other words, they are essentially information-destroying and variety-compressing procedures. As a result, the purely technical approach cannot be very successful. Similarly, the fundamental approach which considers only earnings and dividends in investment analysis cannot yield very good results.

The practical development of cybernetic investment decision systems can be done with the aid of AI techniques. Intelligent behaviors displayed by advance cybernetic systems are problem solving, learning, and pattern recognition. Consequently, I will be concerned with the three corresponding central aspects of AI: heuristic problem-solving methods, learning system theory, and pattern recognition. Since I will not go into the details (2) of implementation of cybernetic decision systems, I will not consider heuristic programming here.

Man/machine work

Because, in the foreseeable future, it will not be possible to automate or perform realistic judgmental problem solving/decision making (PS/DM) situations fully, even with AI techniques, a computer must be used to augment or support rather than replace human judgment.

By augmenting or amplifying human intellect I mean: (1) delegating to the computer processing of those information patterns which are too complex for the human mind to handle; (2) making minds to work better by giving them better tools with which to work; and (3) maximizing the benefits of man-computer "synergism." Although neither the human analyst alone nor the computer alone can consistently outperform the market averages, the combined man-computer decision system may be able to do so.

> **Artificial Intelligence techniques now enable us to efficiently synthesize large amounts of relevant information into investment decisions.**

Learning decision model

Decision making is the process of converting information into action. The investor has a certain amount of information about the general market and some individual stocks and he has a choice among several well-defined investment policies. Typical investment policies include buy, sell, hold, do nothing, sell short, etc. The investment decision selects the best investment policy using available information.

Each investment policy is characterized by a different pattern of information. For example, a stock considered for sale may have a declining earning trend, high price earning ratio, limited growth prospects or pays low dividends. On the other hand, a stock we would like to buy may have a rising earnings trend, good long-term growth potential, relatively low price earnings ratio, innovative management, and a rising price trend on its charts.

Hopefully, there is a one-to-one correspondence between patterns of information and different investment policies. Investment analysis, therefore, involves pattern recognition. Selecting an investment policy corresponds to recognizing the "correct" information pattern. If so, the investment decision process can be programmed or automated with the aid of pattern recognition (PR) techniques.

Weighting evidence

Both investment timing and selection decisions are made by weighing evidence found in patterns of information. Several PR techniques have been developed which are suitable for programming decision-making by weighing evidence. The most useful of them are known as *linear discriminant* techniques (3). Such techniques are computationally simple, need little memory, and can, therefore, be easily installed on personal computers. I recommend these PR schemes become the heart of your programmed investment decision models.

A decision model

A decision model operates as follows: each factor used in investment analysis is first numerically encoded. In other words, investment information like observations of stock market indicators, corporate earnings trends, psychological measurements, etc., must be converted into numbers before they can be processed by computers. We call this ordered set of numbers the *feature vector.* The feature vector represents a numerical encoding of all investment information captured for the decision model.

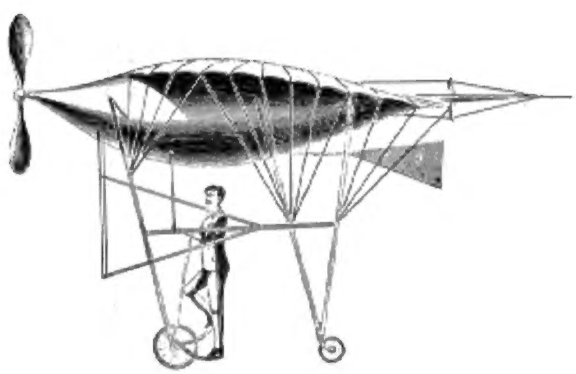

Next each feature is assigned a numerical weight. This ordered set of weights is called the *weight vector*.

Remember that every investment policy has a corresponding unique pattern for which we are looking. Discrimination between patterns is accomplished by means of the weight vectors. To make investment decisions, we form the product of the feature vector with all weight vectors. That is, for each investment policy we form the product of each pattern feature with its corresponding weight and add them up. These sums are the values of the discriminants corresponding to the alternative investment policies. We then select and carry out that policy whose discriminant is largest.

Learning process

Now, the process which generates stock price changes is poorly understood. Therefore we cannot design "good" investment decision systems initially. In other words, we do not know what weights should be

assigned to the features in the pattern recognition mechanism to obtain good performance. Therefore, at the beginning the pattern recognition mechanism will perform poorly and many poor decisions will result.

We can, however, gradually improve performance of the decision system during its operation. This can be done by adjusting the weights of the pattern recognition mechanism so as to reduce the error rate. These weight adjustments are made under the direction of information feedback derived through evaluation of past investment decisions. This process of gradually modifying the decision system so as to improve its performance is called learning.

Perceptrons

Many computer-oriented learning techniques for performance improvement (or learning) of pattern recognition systems have been developed. I recommend a family of techniques known as *perceptrons*. The perceptron algorithm is a model of "reinforcement learning." It will gradually increase the weights of those features that contribute to improved performance and decrease weights of features responsible for deteriorating performance.

The learning process is supervised by some performance evaluation mechanism which uses information regarding the correctness or incorrectness of the decisions made by the system. This information is derived by evaluation of outcomes of investment policies after they were carried out. The supervision of the learning process and system performance evaluation may be performed by the human operator or it may be at least partially automated.

The perceptron scheme is the central information processing element of our man-computer investment decision system. First, it is the computationally simplest PR technique. Secondly, it requires less information for optimal design than other PR methods. Thirdly, it is well suited for programming decision making by weighting evidence. Moreover, the perceptron can process a large amount of decision making information because there is no limit on the number of features it can handle. This is vital because in investment analysis having more relevant information may improve investment decisions.

A decision system must be able to learn and to adapt to changes in the market environment.

In addition, a decision system based on a perceptron scheme is inherently quite reliable in the sense that failures in one or more features may not markedly deteriorate system performance. Much of the information on which investment decisions are based may be "noisy," that is, it may not be very precise or accurate, and it may be subject to various environmental disturbances. Fortunately, the perceptron may continue to function properly even if some piece of input information is missing or incorrect. In this sense our decision model resembles the human brain: it functions in terms of patterns, and if some elements of the pattern are missing or are corrupted with noise, the brain, in effect, reconstructs them.

Finally, perceptron-oriented PR systems need very little memory for computer implementation. The memory needs are small because only the weight vector and the most recent feature vector must be stored. In other words, only the information used of the most recent decisions must be remembered for the learning process. Therefore perceptron-based decision systems can be easily installed on personal computers.

Machine learning

The success of the programmed decision system depends on its ability to predict future direction and changes of trends in stock prices. It is apparently impossible to predict future price changes through an analysis of only past price changes. But if the information pattern for stock market analysis becomes very complex, and includes fundamental data, psychological measurements, news background, in addition to technical factors, then prediction of future price changes may be possible.

This information pattern which serves as a basis for investment decisions is probably so complex that the human mind is generally unable to ascertain the correlation between present values of market indicators and future changes in stock prices. I believe this relationship can be gradually determined through a machine learning algorithm.

Investment selection

The design of every programmed investment decision system must be tailored to the requirements within individual organizations (e.g., investment research departments). The first step in the design process is to specify the user's investment objectives in operational terms, i.e., a selection criterion must be established. Within the decision system these investment objectives are represented by a performance index.

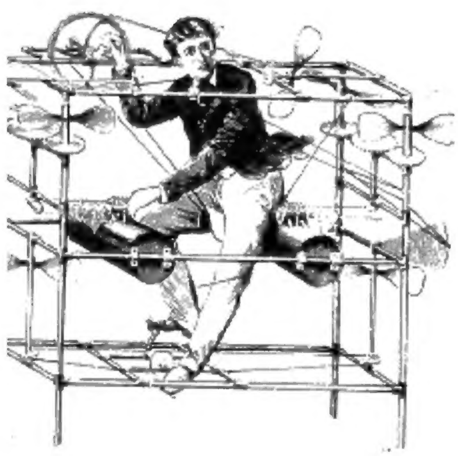

For example, the relative strength criterion is often used. A security is defined as relatively strong over a given time horizon if it appreciates more than a general market index during that time while entailing no more than average (market) risk, and it is relatively weak otherwise.

Clearly, relatively strong issues are characterized by a different information pattern than relatively weak stocks. In other words, this selection criterion divides the set of stocks into two classes: those that are in a relatively strong state and those in a relatively weak state (during a prespecified time horizon). So a two-category pattern classifier can be used to program investment selection according to this criterion.

The second important design step is developing the mechanism for identification of information for decision making in the form of the feature vector. The features may include fundamental characteristics such as earnings trend, earnings stability, future earnings potential, dividend yield, dividend trend, price earnings ratio, risk characteristics (i.e., variance of past price changes or the beta coefficient); technical factors like current price trend on the charts, volume measurements, price overhang, psychological factors (e.g., divergence measurements); and various other factors like news background, capitalization,

quality of institutional sponsorship of the security, company's reputation for research, quality of management, and Standard & Poor's rating of the security.

Next the decision rule will be programmed. The selection rule can be programmed by discriminant functions: select a particular security if the discriminant for that security is greater than zero, and discard it otherwise. This rule will associate a selection policy, e.g., buy (hold) or sell (discard), with any pattern of information describing the security.

Finally, the selection rule is optimized through a learning process with the aim of minimizing the probability of error, as described above. For example, suppose that a security has been selected (bought or held). That is, its discriminant is positive. If the stock subsequently outperforms the market during the prespecified time horizon, the system has made the correct decision. (In this case the weight vector remains unchanged but the counter n is increased by one.) But, if the security appreciates less than the market index, a wrong decision was made, and the weight vector is subsequently changed according to a learning algorithm.

This "correction" reduces the chance of making a similar poor selection in the future. The system is also "corrected" when a reverse error was made, i.e., when a relatively strong stock was not selected. The new value of the weight vector is then used in subsequent computations until the next error is made.

After some time this step-by-step learning from past experiences may result in real performance improvement. At the start, the system's performance will most likely be only mediocre. After about 20 learning steps ($n = 20$), the system may become close to optimal and may make better decisions than human analysts could do intuitively. For example, I am currently using these methods for short-term general market timing: If I get a strong signal, then this approach enables me to predict the future direction of stock trends with about 80 percent accuracy.

Learning about AI

I have provided a broad general overview of the technique. I encourage you to think of the article in just these terms—an overview that establishes terminology, provides examples, and maps the landscape. Entire books have been written describing AI systems that I have described in few words. Each topic discussed has an extensive literature.

I monitor journals and magazines in the field. A good general overview is provided by *AI Magazine*, which is the official journal of the American Association for Artificial Intelligence (AAAI). The magazine is a benefit of membership in AAAI. *AI Magazine* is a blend of general-interest and technical articles that can generally be read by technicians and nontechnicians alike. The magazine also provides information about various meetings and conventions.

The annual convention of the AAAI is a good place to go if you want to see the latest hardware and software systems demonstrated. Tutorials given in conjunction with the AAAI convention offer an excellent way to learn more about the field, often directly from the creators. Membership information can be obtained by writing to the AAAI, 445 Burgess Drive, Menlo Park, CA 94025 (415/328-3123).

With membership in AAAI you are also offered the opportunity to subscribe to *ARTIFICIAL INTELLIGENCE, An International Journal*. This is a scholarly academic research journal, but is worthwhile if your interests are serious. The journal is available from Elsevier Science Publishers, P.O. Box 1991, 1000 BZ Amsterdam, The Netherlands.

William Kaufman Inc., publishes books on AI and sells the conference proceedings of the last several years of both the AAAI and the biennial IJCAL (the International Joint Conference on AI, held every odd-numbered year). William Kaufman also publishes *The Handbook of Artificial Intelligence* (a three-volume reference on the theory of the entire field of AI) and other books related to AI. If your company is setting up an AI group, these publications might be useful. They can be obtained by writing William Kaufman Inc., 95 First Street, Los Altos, CA 94022.

If you are setting up an AI group, you might also consider subscribing to a report service offered by Scientific Data Link, 805 Third Avenue, New York, NY 10022. This company reproduces the reports of the major AI research labs, including Stanford, MIT and Carnegie-Mellon.

At least 15 newsletters have recently appeared that monitor the AI industry and report about new AI research and other developments. The *SIGART Newsletter*, which reports recent AI research results, is published quarterly by the Special Interest Group on AI of the Association for Computing Machinery. The *Applied Artificial Intelligence Reporter* is published monthly by the Intelligent Computer Systems Research Institute of the University of Miami.

There are also several periodicals on knowledge engineering. A useful trade magazine is *Expert Systems: The International Journal of Knowledge Engineering*, published quarterly by Learned Information, Inc., 143 Old Marlton Pike, Medford, NJ 08055.

I estimate that the total number of scholarly journals, trade magazines, and newsletters which are closely relevant to AI, now exceeds 50, and their number is growing fast.

This article is an excerpt from two books by Dr. Jerry Felsen: 1) Cybernetic Approach To Stock Market Analysis, *198 pp., $2; 2)* Your Intelligent Computer Assistant: How You Can Profit From Artificial Intelligence—in Your Business, Home, and Investing, *150 pp., $20. Both books can be ordered from the CDS Publishing Company, 84-13 168th Street, Jamaica, NY 11432, prepaid for $20 per copy.*

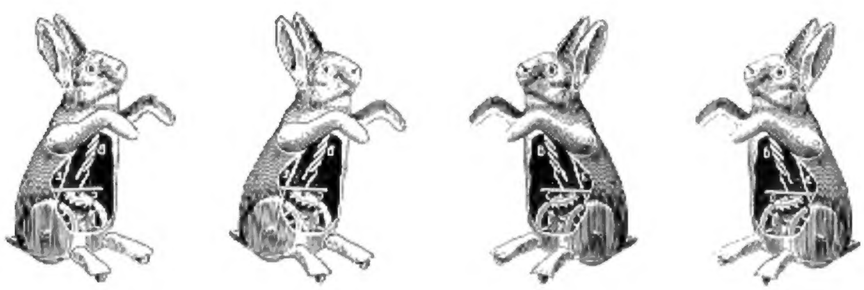

Dr. Jerry Felsen, President of the Center for Applied Artificial Intelligence, is an author, lecturer, computer consultant, and investment adviser. He has an extensive background in artificial intelligence and investment management. He has been involved in investing and development of computer-based systems since 1963. He received his Ph.D. degree in 1973 from the University of Pennsylvania, specializing in computer and information sciences with minors in operations research and investment management. Dr. Felsen has written seven books and many papers on computer technology, investment management, and applied artificial intelligence.

References:

1) *Cybernetic Approach to Stock Market Analysis*
2) *Your Intelligent Computer Assistant: How You Can Profit From Artificial Intelligence—In Your Business, Home, and Investing*

